Towards Addressing the Challenges of Data-Intensive Computing on the Cloud

Eun-Sung Jung and Rajkumar Kettimuthu, *Argonne National Laboratory*

The era of big data is here. The amount and generation rates of data are dramatically growing in science, business, and government sectors and in IT management areas. For example, a tomography beamline at the Advanced Photon Source (APS), an experimental facility at Argonne National Laboratory, can produce 150 terabytes of data in one day. The Data-Driven Urban Design and Analysis project[1] at the University of Chicago aims at utilizing city-level sensor data gathering and predictive analytics on the data for the future urban design. Such data-intensive applications so far have been run on dedicated high performance computing facilities, in order to ensure timely computation for the applications. But this approach can be adopted only by a small number of users who have access to the expensive high performance computing facilities. Even this approach can scale up only to the capacity of the particular facility. Nevertheless, cloud computing can provide any user access to large computing resources for data-intensive applications, on-demand with greater flexibility, at low cost.

In order to deploy data-intensive applications on clouds successfully, however, sophisticated big data management and data transfer frameworks must be seamlessly incorporated into cloud computing infrastructure. In this article, we highlight the current limitations of cloud computing for data-intensive applications, and discuss the challenges facing deployment of data-intensive applications on the cloud and measures to address these challenges.

LIMITATIONS OF CLOUD

Data transfers in and out of the cloud and among different clouds are typically carried out through system utilities such as *Secure Copy* (*scp)* and *http.* These utilities are, however, notoriously slow; indeed, in some cases, shipping disks using commodity transportation services such as FedEx may be faster.

The data throughput of existing data movement services (~tens of megabytes/second) is an order of magnitude less than the throughput required by data-intensive applications (~hundreds of megabytes/second).

Besides data movement among clouds, data management platform within a data center plays key roles. However, popular cloud data management software such as Apache HDFS/HBase, Google FileSystem/Bigtable, and memcached are not directly applicable to data-intensive computing. In particular, because of an unoptimized data movement framework as well as limitations in the underlying hardware infrastructure – shortage of network bandwidth and high latency of data access – cloud computing may not satisfactorily meet the requirements of data-intensive applications. A recent article[2] addresses challenges in migrating data-intensive applications to the cloud from the perspectives of federating diverse data management and hardware systems in the cloud. In this article, we focus more on a data movement framework.

Clearly what is needed are high-throughput and scalable data management techniques as well as higher performance network and storage facilities together with stricter quality-of-service (QoS) requirements on the cloud.

DATA-INTENSIVE CLOUD APPLICATION CHALLENGES

The deployment of data-intensive applications into the cloud presents several key challenges, outlined below.

Distributed data sources. One of the key enablers for data-intensive computing on the cloud is the ability to move data from the data source to the cloud efficiently. Data sources include distributed sensors, genome sequencing centers, experimental facilities such as the APS, and observational facilities such as the Very Large Telescope in Chile. Several challenges are raised by these various and growing data sources.

- The amount of data being generated at data sources is growing because of enhanced sensing capability and the increasing number of data sources.
- Distributed data makes it harder for a single cloud to process the data alone due to longer processing latency and bottleneck in network links to the single cloud.
- Sustaining data transfer rates from data sources to the cloud is difficult because of the dynamic network status and the

interplay of data movement and computation in cloud.

Multicloud environments. Intercloud data movements may happen if a single cloud is not enough for computation or a user or collaborators have multiple cloud resources. While similar to data-source-to-cloud data movements, this kind of data movement is distinctive in that extensive cloud resources can be put for data movement. Two questions arise in this context:

- How many clouds are required to process the data, considering the cost of data movements among them?
- How many resources (i.e., VMs, storages, network bandwidth) in each cloud are required for data movement?

Quality of service. Data center networks are also critical to the scalability and performance of data-intensive applications, requiring a large amount of communication among compute nodes. The usual cloud computing infrastructures in data centers cannot provide scalable data movement performance comparable to that of HPC systems because of insufficient network bandwidth and data movement based on a best-effort policy[3]. Predictable execution of data-intensive applications, presents two challenges.

- The data movement performance of an application should be isolated against other running applications.
- Resource provisioning should be autonomous to adaptively guarantee QoS despite dynamic network status changes.

THE WAY FORWARD

Despite the challenges facing data-intensive applications on the cloud, the forecast is good.

In order to cope with higher data rates from data sources and intercloud movement and to sustain higher data transfer rates, we can exploit two measures: overhead minimization in data transfer and multiple data flow orchestration. Similarly, for better intracloud movement, emerging technologies such as QoS-aware resource management offer opportunities for fine-grained resource provisioning while ensuring quality of service.

Overhead minimization. The current state of the art in high-speed networks is 100 Gbps network links, and several such research networks exist[4]. Approaches for minimizing data transfer overhead include

- eliminating redundant data copies: remote data memory access[5];
- harnessing data locality to avoid delayed data access time: Application CPU allocation aware data flow processing, which allocates application processing and protocol processing on the same core or on the cores that share the cache[6];
- reducing control message wait time, e.g., with GridFTP pipelining for transferring many small files[7]; and
- isolating data transfers from computation, through dedicated transfer nodes (DTNs) and Science DMZ (http://fasterdata.es.net/science-dmz) for a stable data movement framework.

These approaches have not been applied to cloud platforms, however, and some cannot be directly adopted for cloud platforms. For example, in order for cloud architecture to incorporate DTN and Science DMZ, there should be a dynamic resource manager that allocates a proper number of DTNs and reserves storage space on applications' data transfer requests.

Multiple data flows. Multiple data flows already are being used to achieve higher throughput. For example, multipath TCP (MPTCP) exploits multiple data flows (http://www.multipath-tcp.org); but optimized multipath establishment is not a role of MPTCP. Such parallelism can also happen in the disk storage layer, data-processing layer using multicores, network path layer, and application layer using multiple nodes. Settlemyer et al.[8] proposed an impedance-matching technique taking into account multiple layers involved in data movement. Similarly, we proposed graph-based modeling of data transfer systems to optimize data transfer throughput[9] as in Figure 1. Our approach can mathematically formulate the optimization problem based on analytical models of components involving data movement.

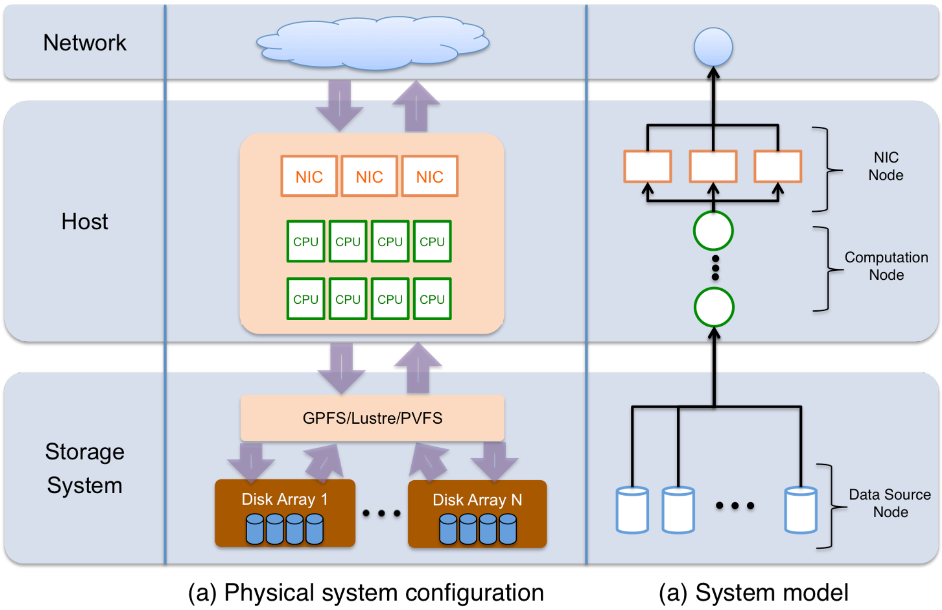

Fig. 1. Graph-based modeling of a data transfer system for data transfer optimization

Nevertheless, improvement is clearly needed in order to accurately model time-varying components. To deal with the inherent dynamic status of wide-area network together with unpredictable performance on cloud resources, researchers should focus on continuous monitoring and building predictable models on time-varying network status and data movement performance in the cloud, in order to guarantee sustainable data transfer rates.

QoS-Aware Resource Management. In addition to data-source to cloud and intercloud movement, intracloud movement also can be improved. Arguably, QoS-aware resource management in the cloud has been difficult because of the unpredictability of application performance on cloud resources and the inability of fine-grained resource provisioning. Nevertheless, emerging technologies such as network virtualization and software-defined network have enabled researchers to do fine-grained traffic engineering and network path control, which may lead to better network utilization for data center networks. We believe the cloud is moving toward advanced operation (or resource) management where all the resources are provisioned in a fine-grained and predictable way, which has not been possible so far. Sophisticated QoS-aware scheduling algorithms should provide holistic resource provisioning while guaranteeing QoS.

A Bright Future

There is growing interest in embracing data-intensive applications, and we envision a bright future for data-intensive computing on the cloud. Realization of this vision, however, requires enhanced data management over intercloud and intracloud networks as well as adoption of emerging network and storage technologies, as we outlined in this article. Some barriers for adoption of such new infrastructures may exist. For example, relatively high cost of advanced network and storage infrastructures such as Dragonfly network topology may deter cloud service providers from deploying them. Nevertheless, the increasing user request for data-intensive applications on the cloud and benefits for cloud service providers such as energy consumption reduction through efficient data movements will accelerate the adoption of new data management infrastructures.

Acknowledgement

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.

REFERENCES

1. "URBAN SCIENCES RESEARCH COORDINATION NETWORK: DATA-DRIVEN URBAN DESIGN AND ANALYSIS." [ONLINE]. AVAILABLE: HTTPS://URBANCCD.ORG/PROJECTS/URBAN-SCIENCES-RESEARCH-COORDINATION-NETWORK-DATA-DRIVEN-URBAN-DESIGN-AND-ANALYSIS.

2. G. CHEN, H. V. JAGADISH, D. JIANG, D. MAIER, B. C. OOI, K.-L. TAN, AND W.-C. TAN, "FEDERATION IN CLOUD DATA MANAGEMENT: CHALLENGES AND OPPORTUNITIES," *IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING*, VOL. 26, NO. 7, PP. 1670–1678, JUL. 2014.

3. A. MARATHE, R. HARRIS, D. K. LOWENTHAL, B. R. DE SUPINSKI, B. ROUNTREE, M. SCHULZ, AND X. YUAN, "A COMPARATIVE STUDY OF HIGH-PERFORMANCE COMPUTING ON THE CLOUD," IN PROCEEDINGS OF THE 22ND INTERNATIONAL SYMPOSIUM ON HIGH-PERFORMANCE PARALLEL AND DISTRIBUTED COMPUTING, NEW YORK, NY, USA, 2013, PP. 239–250.

4. "100G TESTBED." [ONLINE]. AVAILABLE: HTTP://WWW.ES.NET/RANDD/100G-TESTBED/.

5. B. TIERNEY, E. KISSEL, M. SWANY, AND E. POUYOUL, "EFFICIENT DATA TRANSFER PROTOCOLS FOR BIG DATA," IN 2012 IEEE 8TH INTERNATIONAL CONFERENCE ON E-SCIENCE (E-SCIENCE), OCT., PP. 1–9.

6. V. AHUJA, M. FARRENS, AND D. GHOSAL, "CACHE-AWARE AFFINITIZATION ON COMMODITY MULTICORES FOR HIGH-SPEED NETWORK FLOWS," IN PROCEEDINGS OF THE EIGHTH ACM/IEEE SYMPOSIUM ON ARCHITECTURES FOR NETWORKING AND COMMUNICATIONS SYSTEMS, NEW YORK, NY, USA, 2012, PP. 39–48.

7. J. BRESNAHAN, M. LINK, R. KETTIMUTHU, D. FRASER, AND I. FOSTER, "GRIDFTP PIPELINING," TERAGRID, 2007.

8. B. W. SETTLEMYER, J. D. DOBSON, S. W. HODSON, J. A. KUEHN, S. W. POOLE, AND T. M. RUWART, "A TECHNIQUE FOR MOVING LARGE DATA SETS OVER HIGH-PERFORMANCE LONG DISTANCE NETWORKS," IN MASS STORAGE SYSTEMS AND TECHNOLOGIES (MSST), 2011 IEEE 27TH SYMPOSIUM ON, 2011, PP. 1–6.

9. E.-S. JUNG, R. KETTIMUTHU, AND V. VISHWANATH, "TOWARD OPTIMIZING DISK-TO-DISK TRANSFER ON 100G NETWORKS," IN *2013 IEEE INTERNATIONAL CONFERENCE ON ADVANCED NETWORKS AND TELECOMMUNCATIONS SYSTEMS (ANTS)*, 2013, PP. 1–6.

BIOGRAPHIES

Eun-Sung Jung (esjung@mcs.anl.gov) is a postdoctoral researcher in the Mathematics and Computer Science Division at Argonne National Laboratory. He earned a Ph. D. in the Department of Computer and Information Science and Engineering at the University of Florida in 2010. He was a research staff member at Samsung Advanced Institute of Technology from 2011 to 2012. His current research interests include cloud computing, network resource/flow optimization, and real-time embedded systems.

Raj Kettimuthu (kettimut@mcs.anl.gov) is a project leader in the Mathematics and Computer Science Division at Argonne National Laboratory and a fellow at the University of Chicago's Computation Institute. His research interests include transport protocols for high-speed networks, research data management in distributed systems, and the application of distributed computing to problems in science and engineering. He is the technology coordinator for Globus GridFTP, a widely used data movement tool. He has published over 70 articles in parallel, distributed and high-performance computing. He is a senior member of IEEE and ACM.